This chapter and the next are incomplete insofar as they both break off an argument and leave hanging threads to be taken up in the next volume. In this chapter, we review some existing applications of the theory and the derived methods, as well as considering a few novel applications. In contrast, Chapter 11 is very speculative. Some fresh but already promised interpretations are given, but these pose questions that cannot, at the moment, be answered. The answering is left in abeyance until the next volume, though I believe the questions open up an exciting and fascinating field of enquiry.

1. Entailment Structures

It is worth stressing that once the entailment structure for a subject matter is available, together with its task structure, these can be used to design any kind of course or curriculum. With the practically unimportant caveat of Chapter 7 (testing for entry requirements) the structure is a representation of knowables, unqualified, and the derivation of course materials to suit the types of learning which have been described is a matter of (trivial) routine manipulation and (thoroughly non-trivial) literary and/or graphic representation of the topic relations specified. The saving in labour is known to be considerable and is certainly sufficient to place in court the otherwise rather impractical idea of writing different course modules to suit different styles of learning. Provided that a phasing/selection of transmissions is possible the course modules can be television/radio programmes, just as well as books, and if phasing is not available it may be possible, as a second best option, to pack two "streams" of instruction or entertainment (the distinction is questioned in the next chapter) into one presentation.

2. Individual Tuition Materials

Just as the entailment structure can be viewed apart from the construction system, so (to a more limited extent) the tutorial heuristic can be divorced from a facility like CASTE. One exciting possibility, at the moment under scrutiny, is a new kind of publication.

Instead of publishing text books for teachers and students to read let us publish an entailment structure and a task structure, together with tutorial data modules. The entailment structure is no more than a printed sheet together with "magic markers" to be positioned (like the CASTE marker predicates) as tokens of progress, intention and understanding. The task structure usually involves a physical modelling facility and its presence is crucial as a device for eliciting non-verbal and thus easily interpretable explanations. But the facility in question is generally not too costly. As part of the same package, it is possible to provide readily administered tests for strategic disposition and competence; several mini tests (not described in this book) are now known to be passably good discriminators. Apart from the modelling facility (which is shared) each pupil has a set of Individual Tuition Materials.

What could a teacher do, given these "Individual Tuition Materials", which he could not do with a standard text book? It depends upon his inclination and of course, the time at his disposal; the extent to which he can use group instruction and stream the class. But the following possibilities are certainly quite often feasible:

(a) The teacher can use the entailment structure and the tutorial modules (the latter augmented to any desired extent by his own materials) in order to specify a locally tailored course of instruction and investigation.

(b) It is obviously impossible to execute the full uncertainty regulation heuristic manually. However, as noted earlier, the dominant factor in uncertainty regulation is matching between an individually recommended (or student chosen) learning strategy and the student's competence to execute it. Competence can be ascertained by the mini tests mentioned in the last paragraph. A strategy may either be chosen (as a teaching strategy) by the teacher or (as a learning strategy) by the student. In the latter case, advice should be given and a right of objection reserved.

(c) Similarly, it is impracticable to execute the full CET heuristic manually. But it is possible, as a second best alternative, to elicit the crucial explanations that indicate understanding and to demand explanations (obtained through the task structure and modelling facility) for each topic i.e. to obtain a partially non-verbal teachback (as in Chapter 3). Remedial action, if an explanation is not agreed, should (by the theory) consist in eliciting a distinct explanation (having said how the original was mistaken); mere cueing and prompting to patch up a misbegotten essay is insufficient (it would not perform the crucial trick of forcing the student to invent his own way of reconstructing a concept that he understands).

(d) The most intriguing (and also, perhaps, the most practical) use of individual tuition materials is in connection with groups of students, acting more or less independently, or as teams.

By replicating the icons in several dimensions, each conversational system can be represented as an n-person dialogue for $n > 2$. In particular, CASTE has been worked in this mode (as a 3-participant system; 2 students and one CET heuristic) on a more or less equal footing. Preliminary results obtained from this system though not yet reported, are dramatic.

Further, it looks as though the salient features of the n-person (tutorial) conversation may be replicated in a classroom with rather little equipment over and above the "individual tuition" materials.

3. Examination Design

We commented in Chapter 7 that it is possible to exhibit, by a purely mechanical procedure, all the possible which questions, $\underline{\text{PQuest}}^0$, that can be asked about the thesis represented in an entailment structure. Given a modelling facility it is also possible to exhibit all of the explanation demanding questions $\underline{\text{EQuest}}^0$ that correspond to base commands; given further constraints, sufficient to impose the desired limits, it is also possible to generate the $\underline{\text{EQuest}}^0$ corresponding to qualified commands on each topic relation. Finally, there is no difficulty in enumerating all of the $\underline{\text{PQuest}}^1$ that ask about the form of the subject matter within the compass of the stated descriptors.

It follows that if educational materials are derived from an entailment structure/task structure, then all possible modes of examination are available. It is almost a routine matter to mechanise selections that (for example) sample a specified set of topics and ask questions that have not been asked in other examinations. A much more interesting issue is the power of the examination (according to the theory, a cogent examination tests for understanding, in the sense of explaining and explaining how the explanation is reconstructed, but there is no objection to settling for a less ambitious requirement as a compromise between the ideal and something that can be done in the time usually allocated to examining).

4. General uses

Entailment structures have been chiefly discussed in an educational context. But their application is virtually unlimited. For example, structures of this type naturally represent the knowledge in a library; the books, etc, in which it is stored, a description of what may be known and an indexing scheme based upon questions that are asked or can be asked by users.

Historically, it happens that much of the initial work was done in the context of market research, rather than education. This field will receive attention in the next volume and the following comments serve only to fix this application in the reader's mind. For marketing, the entailment structure represents what may be known, either directly or indirectly, about a class of products or facilities. Either the trouble and cost of accessing this data is specified or else the mode of advertising (magazine graphic, television presentation, or printed message). Repertory grid data, in respondent elicited coordinates, may be imposed upon the descriptors of the structure and the content rating given in the field study can be compared with what the respondent would come to know if he actually explored a data base founded on the entailment structure. To elicit exploration strategies in the laboratory (rather than massed data from field interviews) we employed a simulation (the purchaser decision simulation) in which, over a typical period of 2 or 3 hours, respondents culled the information needed to decide upon the purchase of one out of several alternative durables. The respondents exploration strategy, thus recorded, is equivalent to a students learning strategy and it is interesting that respondents showed styles of thinking that replicate, almost exactly, the categories of learning strategy observed amongst students[25]. In this context, a marketing

[25] The substantive data was derived by quantifying exploration/accessing moves with confidence estimates of direct and contingent belief about purchases, obtained by using the BOSS system to sample the pertinent beliefs for each transaction.

strategy (at least to the extent that it aims to persuade or inform or answer questions) is isomorphic with the teaching strategy. Finally, the changes in uncertainty occasioned by reaching a decision, are closely related to the changes observed in learning.

5. Gross Properties of Knowables

One pedestrian but nevertheless valuable use of a system like CASTE is to obtain the vital statistics for a body of subject matter, represented by an entailment structure and an appropriate

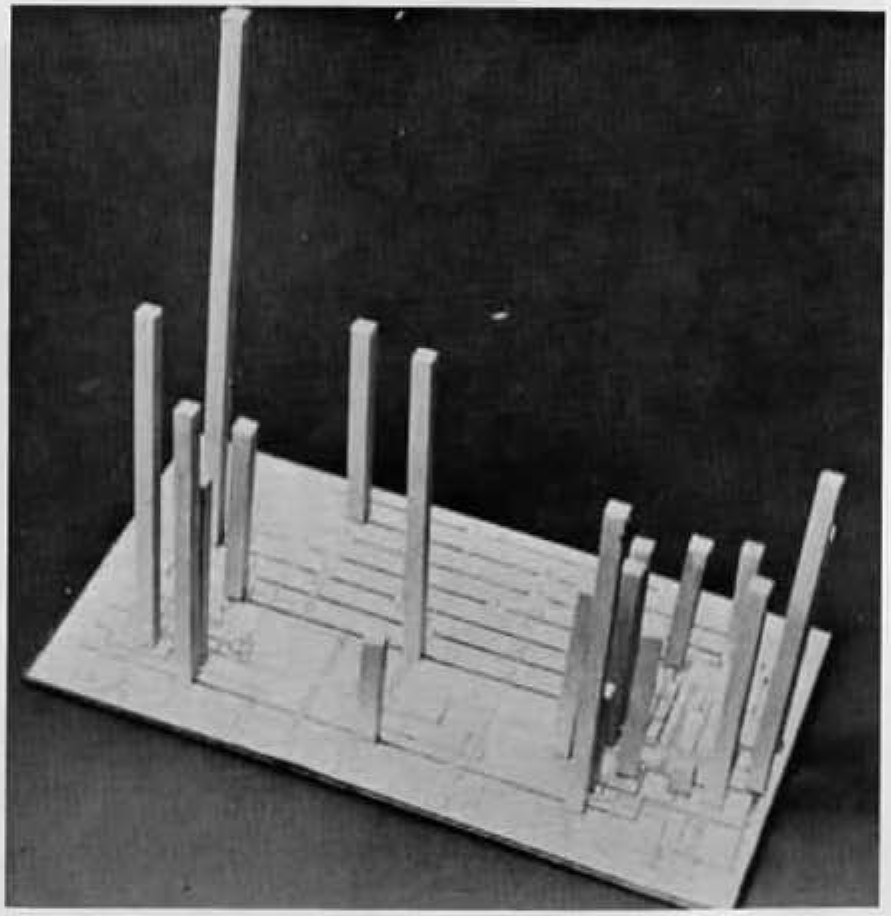

Fig. 10.1. Distribution of mean ΔH per node for the conversational group (half holist, half serialist) (scale values: 0—+2).

task structure, by taking average measures over a sample of students who have used the system conversationally. It is easy to plot distributions, like those in Fig. 1, Fig. 2, and Fig. 3 of (for example) uncertainty (H), correct belief (θ) and the mean interval of study τ. These distributions are rooted upon the entailment structure, regarded as a conversational domain. Such figures are required to set the initial constants and coefficients in an uncertainty regulation heuristic and the need to have them was the chief justification for measuring these values, during the operation of a CET heuristic where they are not used as control variables. The graphs in Fig. 1, Fig. 2, and Fig. 3 represent data from a

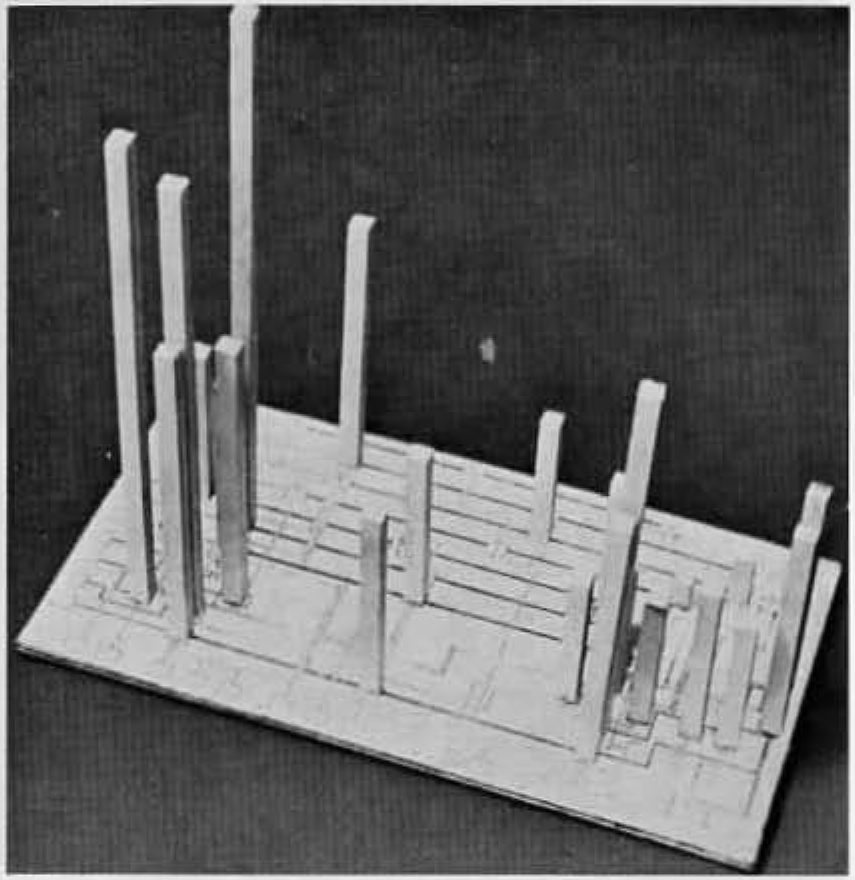

Fig. 10.2. Distribution of mean $\Delta\theta$ per node for the conversational group (half holist, half serialist) (scale values: 0— - 2).

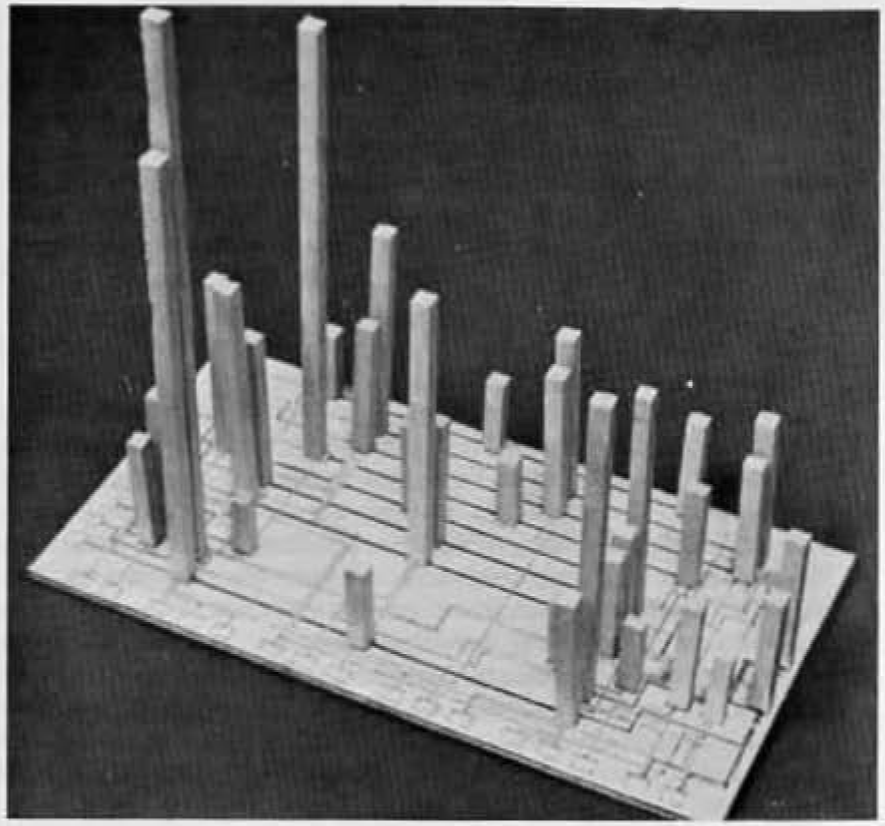

Fig. 10.3. Distribution of mean tutorial time per node for the conversational group (half holist, half serialist) (scale values: 0—20 min.).

deliberately heterogeneous sample of students made up from equinumerous groups of holists and serialists.

Quite apart from the system, such data is of value in general course and curriculum design. It augments the entailment structure by conveying to the designer a good impression of points that give rise to high levels of uncertainty (H peaks), the degree of positive transfer between different topic relations (from the correct belief or θ gradients) and the relative amount of tutorial material that should be prepared to back up each topic.

6. Teaching Teachers

CASTE has been described as a system for constructing course material, for teaching students and for maintaining conversations; namely, as a laboratory facility for studying the learning process. With that in mind, it is appropriate to ask what laboratory facility is needed in order to teach teachers (for example, in courses on educational or industrial psychology; in teacher training colleges, and the like). One answer to this question is, of course, CASTE: the student-teacher being able to use the facility in different roles, and, at the most primitive level, to observe it in operation with real students demonstrating how they learn.

Off all teachers a student may encounter, perhaps the most important, in the long run, is the student himself. We commented in Chapter 3 and Chapter 4 that a pupil and a pedagogue are distinguishable components in the cognitive makeup of any student learning on his own; moreover, that students who have "learned to learn" may reach, in isolation, the standard reached by other people in a tutorial conversation. Conversely, however, there is a positive transfer effect; just as teachers can profit from observing the learning strategies of students (as these are exteriorised in CASTE), so the student who enters the system with this intention becomes aware of his own learning strategies and (in the evolutionary version of Chapter 8) he can master the art of structuring otherwise unstructured subject matter. In other words, CASTE, and probably Individual Tuition Materials also, may be used to teach teachers about learning and also to *teach students how to learn*. Since the aptitudes in question appear to generalise very widely, these applications of CASTE are amongst the most important.

7. General Orientation

The theory put forward has a number of special characteristics. At a microscopic level it is chiefly concerned with strategies (of problem-solving, learning and teaching) and properties of systems of strategies such as compatibility, incompatibility, consistency, inconsistency, reproducibility and cognitive fixity. Since the basic observations are of dialogue in an interpreted language which, in many respects, has the richness of a natural language, the hypotheses proper to the theory can be posed in terms of

conversational modes (explanation, understanding, appreciation and the like) and are tested by observing the occurrence of these modes, the consequences of their occurrence and the configurations in which they occur.

As a further distinguishing mark, the theory is indifferent towards the particular interpretation given to "participants" in a conversation. The unitary P Individual may be a mind inhabiting one brain; a social organisation, inhabiting several brains or several individuals, coexisting in one brain; for example, the case of "learning alone" (which is viewed as a dialogue between distinct P Individuals "learner" and "teacher" executed in the same brain). It would appear that all of these features fit the theory peculiarly well to such complex issues as general decision taking (not to be confused with "decision theory") and the dynamics of education. For the latter application, the theory is also recommended by one of its main postulates; namely that theories of learning have theories of teaching as their duals and vice versa.

Taking a broad view of education (or the psychology/sociology of education) it is customary to carve the subject into a cognitive part (including its epistemological concomitants) and a maturational part (including developmental psychology).

Regarding the cognitive division, the theory maps this part of psychology into a systemic (I hesitate to say mathematical) foundation consisting in the abstract theory of reproducing and evolving procedures; Fuzzy algorithms in Zadeh's (1973) sense, together with concurrently executed procedures; only in the limit onto automata and serial processes. Somewhat too strongly, cognitive division is relegated to a subtheory of P Individuals (all processor constraints taken as given) together with the conversational domains they inhabit.

For lack of space, the maturational division has received no serious attention in this book. The omission will be remedied in the next volume. The following remarks preview the form of theory to be presented.

We might, quite optionally from the present stance, regard maturation in two ways. From one point of view maturation, imprinting, early learning and so on, equip an M Individual, born as a baby, with the mental equipment needed to live and learn; namely, as the P Individuation of an M Individual. This point of view is adopted by most psychologists, though Vygotsky (1962) saw that it was not the only point of view. At the other polar extreme, and just as consistently, it is possible to regard a culture or a civilisation as providing a gaggle of P Individuals which are M Individuated by the babies born as processors.

In either case, of course, the points of view interact when the system is realised. The dynamic entity of interest is (in the strict sense) a hybrid. But it has been possible to circumvent a number of common puzzles and ambiguities by treating maturation in an other-than-traditional way i.e. by taking the maturational division as the subtheory of M Individuation of P Individuals existent or latent in society.

8. Augmenting Brains

A closely related question is "what should be educated". The trite reply, "children (perhaps adults)" is unacceptable as it stands because the word "child" is ambiguous. It could mean a certain kind of M Individual (aged 8 or so) or it could mean a certain kind of P Individual or, in fact, many combinations of these entities.

For administrative purposes, it is usual to think of the objects of education as brains (M Individuals) and to work out various cost effectiveness criteria for when and where they should be educated. However, it is quite useful to break away from the subculture of head counting and to regard the basic units as P Individuals.

Now, if P Individuals are the proper targets of education and the same cost effectiveness criteria are employed, it turns out that (elitist though it sounds) it is only worth educating a certain kind of P Individual; say a P Individual more complex than a Minimal P Individual. In fact, the break point is very clear. The Minimal P Individual can learn to learn, whereas P Individuals who are not up to the minimal standard fail to do so and in practice are destined to a trapping condition of indoctrination or ossification. On even cursory scrutiny of the data, this makes something like a 10 or 20 fold difference in learning efficiency.

Any functioning P Individual is executed in a processor. At the insistence of the administrator, educational science sets up a very dubious one to one correspondence between brains and minds (as M Individuals and P Individuals) and under this identification there appear to be two pathologies. One of these is quite remediable. A person (P Individual and M Individual to execute it) fails to learn because he has not been taught to do so. By common

agreement "learning to learn" is an art that can be taught (our own instantiation of teaching the art by engaging the student in a CET conversation has been discussed but is by no means the only technique).

The other pathology is less easily overcome. It may happen that the M Individual imposes processor constraints such that it cannot execute a minimal P Individual even if this entity is compiled as a procedure. Apart from gene mechanics, the only solution to this problem is to augment the brain in question. Such an augmentation is fairly common though (because of administrative convenience) people rarely admit that they are up to this trick. Perhaps the most dramatic manifestation is seen in institutions for the subnormal where, in fact, the "students" are organisations residing in groups of brains, which, taken together, can execute the procedures of a minimal P Individual, provided, of course, that there is an appropriate environment of other brains (as there is inside the institution). With this example in mind, it is easy to cite others; the man who is only motivated to learn in company, the whole gamut of familial support, and so on. In each case the teachers responsible would probably agree that the target of education is a P Individual only capable of execution in a cluster of brains.

To date, nearly all augmentation tricks (whether they make it possible to execute a minimal P Individual or enhance the capabilities of a P Individual above this standard) have been played by adding brains to brains. Certain obvious difficulties beset the endeavour. Brains are bulky and quite expensive to transport so that the P Individual looses mobility; the ecological niches available are often unsuited to maintaining a sufficient group of human beings, and so on. Perhaps the most important contribution of technology to education is that it is now possible (and in five years or so it will be commercially feasible) to furnish computing aids that are portable enough to be permanently attached to a human being as a prosthetic for augmenting his brain (and it may be noted, parenthetically, that brains have some outstanding defects as processors). The form and operation of such devices will be discussed in the next volume.

In summary, the hoary but seldom mentioned augmentation trick can be done in a different way, without necessarily forming a cumbersome colligation of brains. The target of education is "child plus augmenting device", such that the "combined processor" is able to execute a Minimal P Individual.

Moreover, such targets of education are cost effectively educable. It is surely worth some trouble to secure the necessary preconditions for ethical reasons or aesthetic reasons (mental cripples are remarkably ugly). Recent data on equipment costs indicate that the exercise can be justified, as well, on purely economic grounds.

Chapter 11. Consciousness, Self-Reference, and Some Tentative Relations between the Microtheory and the Macrotheory

The dominant theme, obtrusive amongst a potpourri of topics considered in this chapter, is consciousness and its relation to macrotheoretic indices of uncertainty (H), correct belief (θ) and the like. Since there is a tendency to avoid the subject of consciousness as ineluctable, or even because it is held in disrepute, many familiar examples are cited and, as a result of that, the presentation is longwinded and slightly redundant.

Although much of the argument is relegated to the next volume, states of consciousness have so much pragmatic importance (especially in education) that it is necessary to provide an overview of our position in this matter, if the reader is to be left with even a partly completed story (in which, for example, the microtheory and the macrotheory are rationally connected though not at this stage fully unified).

1. Consciousness and Synchronicity

The occurrence of understanding in a strict conversation between A and B, localised in L processors α and β, implies a synchronisation of procedures executed in α and others undergoing execution in β. It is "obvious" (in a sense) that the execution of these procedures must be synchronised in order that the transactions of understanding shall occur; it is, perhaps, less obvious that in the absence of the systemic form of an understanding, as it is shown in Icon 3 and Icon 5, there is no reason to suppose that procedures executed in α are synchronised with those in β and that if this synchronisation occurs because of the realisation of this systemic form, then there is a transfer of information, i.e. the understanding "is the case" (not just "a simulation of understanding is performed"); or, to put the matter strongly enough to invite criticism, A and B are, under these circumstances, aware.

In fact, if it happens (as it does in the icons) that A and B are separated (in α and β) by an interface and if the sprout of the A,